End to End Performance Testing Initiative Analysis - DRAFT

Created: *2000-12-05*
Last Updated: *2001-01-03*
C. de Luna

Problem Statement:
While an infrastructure to provide advanced IP based data transport services is largely in place in the Internet2 community, there currently does not exist a well defined, always-on, systematic, and automated approach to characterizing the quality of service parameters of all the components involved in data transport services from a source to a destination.

Goals:

From the general project plan of the Internet2 End-to-End Performance Initiative, this initiative calls for the creation of:

- a persistent, proactive, and widely deployed performance measurement infrastructure, including tools and instruments for detection and resolution
- distributed, coordinated "Performance Evaluation and Response Teams" (PERTs), information resources, and mechanisms for access to expertise
- ongoing outreach, technology transfer and dissemination of best practices to the Internet2 membership and beyond

This discussion will focus on the first goal.

Elaborating on the first goal, this initiative should provide a scalable and extensible test and measurement framework to quantify the quality of data transport service. This framework should allow network engineers to map a path, identify relevant points, and obtain quantitative values from every relevant point in the path. These capabilities will assist in the identification of path and component characteristics and measurements that can be used to ascertain the quality of each relevant point and ultimately the entire end-to-end path. In this discussion a relevant point includes a network interface, the equipment supporting that interface, and links attached to that interface. This framework would scale and facilitate the addition of new features and enhancements. While the focus of the framework must be to support current applications, thought should be given to future extensibility (i.e. end to end measurement support for QoS, IPV6, and other emerging capabilities).

Demarcations

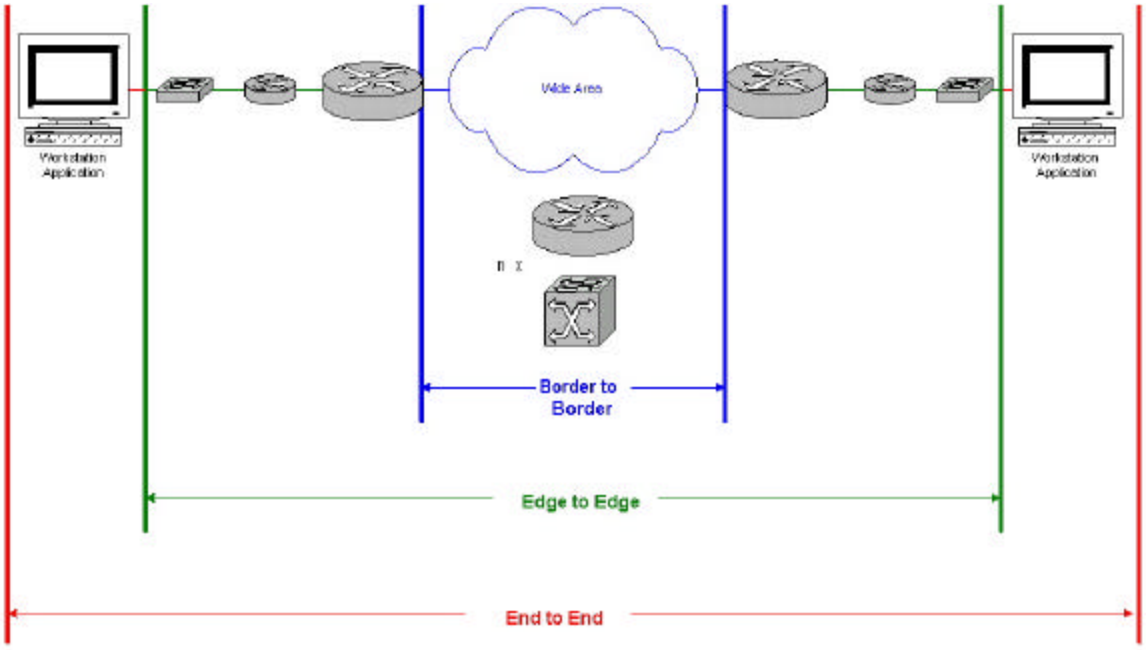

Figure 1 Demarcations

While the focus of this effort is end to end, there may be value in defining segments or demarcations of interest. This will help scope tasks and focus efforts. It also may identify characteristics found in different domains as well as provide a fall back plan should there be issues (political or technical) in deploying instrumentation down to the end host. This type of demarcation can provide well-defined measurement domains that can be used in such cases.

In this discussion 3 fundamental demarcations will be considered. The table below defines them, begins to identify any current capabilities in place to measure characteristics within the domain, and begins to identify any issues inherent in the particular demarcation. This table is by no means comprehensive.

Demarcations	*Description*	*Current Capabilities*	*Issues*
Border to Border	The testing framework will define the characteristics of the path from the border of the source autonomous system to the border of the destination autonomous system.	AMP Surveyor	Crossing service provider organizational boundaries
Edge to Edge	The testing framework will define characteristics of the path from the edge test point	Memory to memory transfer tools	Crossing campus organizational boundaries

	closest to the source (wall jack, telecomm room) to the edge test point closes to the destination.	"Provides High Water Mark" – K. Walsh NetIQ	
End to End	The testing framework will define characteristics of the path from the source to the destination, including the actual systems and applications.	Application	Crossing campus organizational boundaries

In fact, each demarcation may be self-similar. The functional characteristics of the relevant points within the Border to Border demarcation may be very similar to those included in the Edge to Edge. Technology may differ but function will be similar, if not the same. On the other hand these demarcations should serve to identify any differences within the defined demarcations which may impact performance and which should be more closely scrutinized and understood. Demarcations may serve to allow measurement solutions to scale. Demarcations will also be significant in that some organizations may have issues with external entities deploying instrumentation to their edge and to their hosts. Any instrumentation developed or utilized for this effort must take security and monitoring capabilities by the "hosting" organization or campus into consideration. Ultimately, cost and "intrusiveness" of any instrumentation must also be considered. The demarcations defined above are merely examples.

Functional Component Breakdown

As with any effort to solve a problem, a successful solution will depend on the level of understanding of the problem and by how well all the functional components are understood. This is in fact one of the reasons end-to-end troubleshooting is so difficult. Not only is it difficult to measure variables of interest, often it is not clear which variables are of interest or what a value for that variable might mean. There is often no reference data for any values that are obtained. The solution must begin with a clear understanding of the problem, the identification of all functional components that have an impact, and an understanding of how those functional components work and fail. Without this understanding and clarity, scope definition and prioritization may be impossible or inaccurate and solutions may be ineffective.

One possible method, which can be used in this breakdown, is the OSI Reference Model (groan!). Whatever the method, this breakdown is key to determining areas of focus and success criteria.

This breakdown approach needs to be applied overall and in fact to each functional component in the path. In the case of the OSI reference model, not all "layers" may apply to the host, layer 2 Network Device, or layer 3 Network Device. Ultimately, functional components can be grouped into classes with each class having a set of variables of interest. The intent here is to look at the makeup of each functional component, understand how it achieves its function, and understand its failure modes. From this understanding, the necessary variables which must be measured can be identified. Once they are identified, an analysis of current measurement capabilities can be made. In some cases, measurement tools and infrastructure may already exist. These tools or infrastructure may be integrated into the effort, or copied, modified, and/or improved as appropriate. This exercise will identify infrastructure that is non-existent and allow the effort to focus on all or the most important of these. This knowledge will be key in establishing partnerships with organizations, including vendors, to develop the necessary but missing infrastructure.

The table below is intended to illustrate the possible use of the OSI model to look at variables of interest and to begin listing tools and other measurement capabilities for each layer. The last 3 columns attempt to apply this layered approach to other functional components such as the end host, layer 2, and layer 3 functional components which will be encountered along an end to end path. For example, the source and destination hosts must be evaluated in terms of the application being used, how that application interfaces with the lower layers including the operating system (OS), how the OS interfaces with the IP stack and the physical system components (memory, CPU, bus, disk, NIC). The

intermediate networking devices must be evaluated at a minimum in terms of their packet forwarding capabilities.

Layer	*Values of Interest*	*Tools*	*Host*	*L2 Network Devices*	*L3 Network Devices*
7--Application	Application characteristics	?	Application Characteristics		
6--Presentation					
5--Session	OS	?			
4--Transport	TCP/IP Stack/OS	?			
3--Network	Path, dropped packets, queues/buffer measurements Impacts of going through a L2 device *Note: Need to define for vendors*	Ping, traceroute ,pathchar			
2--Data Link	Discarded Frames Runts, CRC, impacts of going through a L2 device *Note: Need to define for vendors*	Interface statistics Monitoring Tools	NIC BUS		
1--Physical	Loss (db), NEXT *Note: How do problems with these values manifest themselves in the upper layers. Can problems here be derived from upper layer data?*	OTDR, cable tester	Memory BUS CPU Disk	Buffers Interfaces Backplane	Buffers Interfaces Backplane CPU

For each demarcation and for each variable, the following capabilities are essential:

- Baseline
  Baseline characteristics for a path and for a variable must be obtained. Inherent in this baseline is the understanding and recognition of the "know working state" from which any deviation may indicate a problem.
- Interactive
  The ability to obtain values for variables of interest in real time at any given moment is desirable. Real time values can be compared with the baseline to assess the quality of a path or component.
- Trending

The ability to keep a history of these values is key to identifying trends in usage and quality. This capability is key for capacity planning as well as for identifying service degradation or variation from the baseline.

Both active and passive measurement abilities must also be taken into consideration.

1) What are the key aspects of end-to-end performance that should be addressed in the Initiative?

For this initiative to be successful the answer to this question is key and resources should be allocated to this breakdown of the functional components along a path so that quantitative answer can be derived. These components include applications, hosts, network devices, and network connections.

Because the data network provides IP based connectivity, the initiative must focus on TCP and UDP at a minimum. Key applications should be identified and characterized. A thorough understanding of TCP and UPD and the performance impacts on each protocol of the behaviors listed below must be considered and prioritized. The table below lists characteristics and behaviors which have an impact on performance and starts to list any existing measurement capability.

Behavior	*Existing Measurement Capability*
Loss	
Duplication	
Out of sequence	
Checksum Failure	Interface statistics
Fragmentation	
Application dependency for error correction for UDP	Application specific
1 way delay	ping
Latency	ping (sort of)
Jitter	iperf (sort of)
Throughput	ftp, iperf, ttcp, nttcp
Tuning	
Packet colors (QoS)	
Path and path changes (routing)	

Once this is determined, measurement criteria can be derived and focus can be placed on the instrumentation necessary to make these measurements.

For example, packet loss has a severe impact on both TCP and UDP and therefore should be at the top of the list in terms of providing instrumentation to identify where packet loss is occurring along an end to end path and why.

2) What are the proper success metrics for this Initiative?

The identification of key applications
The identification of key factors (i.e. behaviors) which impact application performance

The design and implementation of a measurement capability for measurement of those key factors

The correlation of these measurements to provide a meaningful end-to-end quantitative value for the quality of the data service

The widespread dissemination of this framework within the Internet 2 community at a minimum

The successful deployment of PERT teams to assist

The generation of “best practices” and “cookbook” documentation to assist sites in this effort.

Metrics:

- Number of key factors that can be measured end to end
- Number of failure modes that can easily identified in an end-to-end path
- Number of key factor quality thresholds defined (i.e. lookup table defining good, bad, and marginal)
- Number of applications which for which baseline and trending data and real time measurements can be obtained using the measurement framework
- Number of operational sites that support the measurement framework
  % of e2e sites
  % of Internet2 sites

Later gears:

- Number of PERT Teams available for “consulting and engineering”
- Number of PERT Team calls
- Number of “Best Practices” Papers

3) What criteria should be used in the call for participation?

Desirable Site Characteristics:

- Resource commitment
  - End to End work
  - Local Campus outreach
  - Training and mentoring of local NOC staff
  - PERT membership
  - Documentation
- Knowledgeable and committed personnel

- Existing commitment to measurement and monitoring
- Completed or planned advanced network infrastructure
- Agreement to deploy instrumentation to the edge and to the end host

While the inclination may be to select sites with advanced network capabilities, this may actually be a dis-service to the effort. In fact some sites with less desirable network infrastructures may be desirable. A resource commitment from a potential site with knowledgeable personnel as primary contacts for the effort may be more valuable than a site with an advanced network infrastructure. A site that has a proven commitment to measurement, regardless of the state of the campus network may also be more desirable. Familiarity with measurement and monitoring instrumentation will be a key skill set for this effort.

Notes:

Define end to end

Define time synchronization requirements

Define interfaces and their requirements

Define components of an end-to-end link

Define relevant test points

Define the characteristics of the framework
    Different for each path

Questions:

Other than manpower, component installation, space, what other resources will a site need to provide?? PERT Team membership, local training, collaboration on best practices papers.

Once we define the measurement parameters of interest and prioritize them do we want to focus on application characteristics?

What, if any, are the plans for vendor participation and partnering and evaluation of cots instrumentation?

What are the failure modes that impact performance?
Which are the most important? What are the characteristics of those failure modes?

Decompose the problem:

What are the key sets of protocols/applications?

What factors impact the performance of these key protocols/applications?

How can these factors be measured and are there any existing capabilities for measuring them? If so what are their strengths and weaknesses. If you could rebuild them what design decision would you make? If you could add features, what would you add?

Where can/should these factors be measured? Are there any issues with measurement at these points? How can quality parameters for these factors be measured along a path and correlated? Are there any social issued to deal with?

What are the normal values for these factors and the abnormal values indicating a problem? What are the thresholds of interest?